

BeSound: Bluetooth-Based Position Estimation Enhancing with Cross-Modality Distillation

Hymalai Bello
DFKI and RPTU

Kaiserslautern, Germany

Sungho Suh
DFKI and RPTU

Kaiserslautern, Germany

Bo Zhou
DFKI and RPTU

Kaiserslautern, Germany

Paul Lukowicz
DFKI and RPTU

Kaiserslautern, Germany

Abstract—Smart factories leverage advanced technologies to optimize manufacturing processes and enhance efficiency. Implementing worker tracking systems, primarily through camera-based methods, ensures accurate monitoring. However, concerns about worker privacy and technology protection make it necessary to explore alternative approaches. We propose a non-visual, scalable solution using Bluetooth Low Energy (BLE) and ultrasound coordinates. BLE position estimation offers a very low-power and cost-effective solution, as the technology is available on smartphones and is scalable due to the large number of smartphone users, facilitating worker localization and safety protocol transmission. Ultrasound signals provide faster response times and higher accuracy but require custom hardware, increasing costs. To combine the benefits of both modalities, we employ knowledge distillation (KD) from ultrasound signals to BLE RSSI data. Once the student model is trained, the model only takes as inputs the BLE-RSSI data for inference, retaining the advantages of ubiquity and low cost of BLE RSSI. We tested our approach using data from an experiment with twelve participants in a smart factory test bed environment. We obtained an increase of 11.79% in the F1-score compared to the baseline (target model without KD and trained with BLE-RSSI data only).

Index Terms—Smart Factory, Multimodal Fusion, Knowledge Distillation, BLE Positioning, Ultrasound Localization.

I. INTRODUCTION

Nowadays, factories have been transformed into smart workplaces with automation at various levels. Using technologies to optimize manufacturing processes and improve overall efficiency. And focusing on human-centered artificial intelligence rather than technology-driven implementations to ensure a safe and inspiring environment for workers. Hence, human activity recognition and positioning systems are being developed as part of the framework to enable continuous monitoring of human behaviors. Tracking the workers' position is a relevant safety task. For example, on a petrochemical factory floor, chemical asphyxiation can provoke unconsciousness or death through suffocation. Some chemicals vaporize and mix with the air without showing any warning properties that raise the risk of oxygen deficiency [1]. Thus, knowing the worker's position will increase the chances of the worker's survival by speeding up the arrival of the rescue team.

For indoor personnel localization, vision-based solutions are the most accurate [2]. For example, in [3], a fusion between a vision-based, and radio-based real-time location systems (RTLS) for person identification is presented. The authors simulated the scenarios with images generated by a 3D game engine. And they used YOLOv3 for the object detection model combined with RSSI-based localization. They obtained a representative Rank-1 accuracy of 88.7% with 4 people for

2 seconds window. A review of recent advances in vision-based indoor navigation is available in [4]. However, concerns about worker privacy and technology protection are relevant aspects to be considered. This led to alternative approaches. The state-of-the-art non-camera-based methods fall into four categories: dead reckoning technology, optical methods, radio-frequency-based (RF) methods, and sound-based localization systems.

Dead Reckoning is based on tracking the user/asset by estimating the current position in a cumulative manner, starting from a known position. It is not dependent on external infrastructure, is low cost, has low latency, and can provide real-time tracking as it continuously updates the position based on data from previous movements. Due to cumulative prediction, its error also accumulates over time. This means that small errors in the measurement can lead to significant inaccuracies in the estimated future positions. It is also sensitive to sensor calibration and initialization error of the starting known position. Hence, this technology is fused with other modalities in the literature [5]–[7].

Optical Methods, usually emit ultra-short light pulses and measure the time of return to the source [8], [9]. These methods use time-accurate equipment and high-power laser beams that could exceed the eye safety limit [10]. The main drawback of these methods is their sensitivity to environmental distribution and lighting conditions.

Radio Signals and light waves are electromagnetic in origin and do not need a medium to propagate. They differ in their radiation frequency (energy) and wavelength. Radio-frequency signal has wavelengths at least 5000 times longer than visible light. This means that the RF signals experienced less diffusion when they encountered an obstacle/wall compared to light pulses [11]. The authors in [11], used a radar-based solution to provide accurate around-corner indoor localization. The motivation is that it is pertinent to ensure that unmanned robots do not collide with factory personnel who suddenly appear around a corner. Ultra-wideband (UWB) is an accurate RF-based tracking solution. Its limitation is mainly based on cost and the fact that it requires dedicated signal receivers. Moreover, RF technology such as WiFi and Bluetooth is highly scalable and low-cost. WiFi and BLE-based solutions are cost-effective, as the technology is available on smartphones and is scalable due to the large number of smartphone users. The accuracy of such solutions is usually in meters. WiFi signals are more stable than BLE signals [12]. But, with proper signal processing, the BLE method can achieve results comparable

WiFi method and require six times less power than WiFi signals [12]–[14]. On the other hand, BLE-tracking is sensitive to overcrowdedness [15]. The existence of human bodies is proven to influence Bluetooth signals. This can be used as an advantage if the purpose is to monitor crowds.

Ultrasonic Waves are mechanical waves, requiring a medium to travel. The accuracy of ultrasonic sensors has proven to be in centimeters [16]. In addition, they suffer less interference from obstacles and walls than WiFi/Bluetooth technology. They bounce off walls/objects unless the transducers are attached to the object [17]. This makes ultrasonic location systems reliable and less prone to interruptions. However, it needs the installation of ultrasound transmitters and receivers in the environment (customized hardware), increasing the cost.

In this work, we propose to enhance the low-cost and scalable option of BLE-RSSI-based localization with ultrasound coordinates. We employ the teacher-to-student scheme. The teacher-to-student scheme is known as knowledge distillation [18]. The idea is to improve a smaller target model with a cumbersome teacher model’s guidance. In our work, the teacher fused RSSI signals with ultrasound coordinates to leverage multimodal and multipositional information. Then, the student (target model) is guided by the teacher. Once the student is trained, the position of the user is inferred by a model whose inputs are only RSSI channels. The student is maintained only with RSSI signal dependency, which reduces cost, improves performance, and increases the ubiquity and scalability of the solution. The **main contributions** of this work are described as follows:

- We present a solution based on knowledge distillation to enhance the low-cost and scalable localization method based on BLE-RSSI signals. The multimodal and multipositional teacher fuses BLE-RSSI signals with ultrasound coordinates and guides the student. Once the student is trained, the user’s position is inferred with only RSSI signals as inputs for full cross-modality transfer learning.
- We have conducted an experiment to track the position of twelve participants in a smart factory test bed. In our approach, the BLE transmitter only needs to be enabled, and the receivers are distributed around the factory modules (see Fig. 1). Hence, no further application is deployed on the user’s side. The ultrasound data is collected at the same time as the BLE data for synchronized and fair data collection.
- We have evaluated the approach with three student models. The students (target models) were designed with the same neural network structure but with a different number of parameters. Experimental results show that with the smallest student with 754 parameters, the F1-score is 8.98% higher than the target model baseline. Moreover, for the biggest student with 4594 parameters, the F1-score is 70.96%. This represents a difference of 8.89% in the F1-score compared to the teacher with 20416 parameters, and it is 11.79% higher than the target model baseline. These results were achieved by training

the distilled students for 10 epochs. In contrast to the target model’s baselines trained for 100 epochs.

The paper is organized as follows. Section II, give an overview of the background and related work in BLE-based localization methods, ultrasound-based localization methods, and multimodal knowledge distillation. Section III, presents the hardware and experimental setting, and the BeSound knowledge distillation methodology. In Section IV are the results and discussions, including a list of limitations with an overview of future directions. Finally, Section V concludes the work.

II. BACKGROUND AND RELATED WORK

In this section, the motivation for the selected approach is expanded further and a background about the technology fused is given. It is divided into three subsections; Bluetooth Low Energy-Based Indoor Localization Methods, Ultrasound-Based Indoor Localization Methods, and Multimodal Knowledge Distillation.

A. Bluetooth Low Energy-Based Indoor Localization Methods

The BLE positioning system involves two elements; an anchor(s) and a tag(s). Usually, BLE anchors also called beacons are positioned at fixed and known locations as emitters/receivers. The main idea is to estimate the tag(s) position, which is the emitter/receiver. For an anchor-based localization system, several quantities can be measured to estimate the tag position. The literature presents approaches based on the Angle of Arrival (AoA), Time of Flight (ToF), and Received Signal Strength Indicator (RSSI) [19], [20].

The Angle of Arrival (AoA) technology has been introduced to Bluetooth Low Energy (BLE) since Bluetooth 5.1 Core Specification to add directionality to BLE localization systems [21], [22]. The AoA method intends to increase the location accuracy compared to the received signal strength indication (RSSI) approach. With the adoption of AoA, BLE localization systems can provide real-time indoor location, achieving an accuracy of 0.1 meters, which is 10 times higher than traditional indoor positioning technology (RSSI-based). Commercially available options for this technology are BlueIOT (Beijing) Technology Co., Ltd ¹, and Dusun IoT ². For an accuracy of less or equal to 0.1 meters, the technology requires customized hardware deployed in the environment as anchors and real-time location system (RTLS) tags [23]. Then the anchor with multiple antennas in the receiver calculates the AoA of the signal transmitted by the RTLS tag(s) using an RTLS location library. The RTLS library is in charge of processing the in-phase quadrature (IQ) samples received from the Bluetooth stack and implements multipath detection and azimuth and elevation calculation.

The ToF method provides accuracy in the centimeter order [24], [25]. The idea is to calculate the time taken by the electromagnetic wave to propagate from the transmitter to the

¹<https://www.blueiot.com/>

²<https://www.dusuniot.com/resources/technical-brief/the-benefits-of-ble-aao-and-how-you-can-use-it-in-life/>




Fig. 1. Experimental Setting Map. **Left** Volunteer Wearing the BLE-RSSI Emitter (Upper-Arm) and the Ultrasound Hedge (Lower back **BLACK BOX**) Together with the Spatial Position of The RSSI-Receiver in Green Color and the Position of the Ultrasound Beacons in Red Color. **Right** Top View Showing the Distances of the RSSI-Receiver and Ultrasound Beacon in the Smart Factory Test Bed.

receiver, which is proportional to the distance between them. The main conditions are; knowing the positions beforehand of at least three anchors for 2-dimensional tracking (4 anchors for 3D), a high-speed clock, and a common clock reference to be shared between the receiver and transmitter. For example, the propagation speed of an RF signal is around 300000 km/s. This means that to achieve a tracking accuracy of 30 cm the BLE radio and the timer employed for measuring the ToF should be clocked at 1 GHz.

On the other hand, the **RSSI-based** localization technique is low-power, affordable, and simple to implement with standard hardware available in traditional versions of BLE [19]. Even though the accuracy is impacted by the line of sight, the system still captures an RSSI tag in a limited way. Moreover, it possesses an accuracy of more than one meter and is sensitive to crowded places [15]. But still is considered a sufficient method for tracking personnel inside a building. Hence, different solutions have been proposed in the literature to improve the RSSI localization systems. The solutions include multichannel transmission using RSSI signal aggregation techniques, reducing the positioning error from 1.5 m to about 1 m. The authors in [26] employ a Kalman filter to fuse RSSI and Time-of-Flight measurement data, and demonstrate the benefit of not relying only on RSSI, comparing ranging performed with and without the help of ToF. In [27], a fusion between AoA and RSSI with a two-step approach is introduced. First, the AoA measurement is smoothed using a Kalman filter. Then a convolutional neural network (CNN) is used for feature extraction of RSSI and AoA, respectively. The fusion is performed at the feature level by a concatenating operation. Moreover, in [28], the authors present a deep neural network-based method that uses the RSSI and the in-phase quadrature (IQ) components of the BLE signals received at multiple anchors to estimate the AoA. The AoA is then used to achieve a localization accuracy of 70 cm.

Overall, the BLE and RSSI-based localization method offers reduced complexity, low power, and scalability. It is a cost-effective solution, as the technology is available on smart-

phones and is scalable due to the large number of smartphone users. Moreover, in our approach the movable tag (e.g. user's smartphone) needs only to be enabled, no further application is deployed on the user's side. The receivers, which are located on the factory modules capture the RSSI signal from the tag without even being paired to the BLE client.

B. Ultrasound-Based Indoor Localization Methods

Ultrasound-based tracking has two elements; a hedge and the beacons. The hedge(s) is then tracked by the ultrasound beacons. Where the beacons' position is previously known. Ultrasound RTLS is often used as ground truth in the literature [29], [30]. It provides high accuracy in location tracking, often achieving centimeter-level precision. The ultrasound wave propagates at 340 m/s in the air which is almost 1 million times slower than the BLE signal (RF wave). In theory, this indicates that it is 1,000,000 easier to achieve higher accuracy with ultrasound-based RTLS than with radio waves, such as ultra-wideband (UWB) based RTLS. Compared to WiFi/Bluetooth technologies, ultrasound technology typically experiences lower interference from obstacles, walls, or other electronic devices. This makes ultrasound RTLS reliable and less prone to disruptions. However, compared to WiFi/Bluetooth-based localization, the ultrasound method requires additional infrastructure costs. It needs the installation of ultrasound transmitters and receivers. The need for a higher density of transmitters in certain environments will increase the complexity and cost of the entire system. All of the above indicates that ultrasonic system accuracy (centimeter-level) is desirable in RTLS. Still, the simplicity, cost-effectiveness, and scalability of a solution such as BLE-RSSI localization are also desirable in a positioning system.

C. Multimodal Knowledge Distillation

The large-scale deep neural network models have been successful in many fields, achieving high performance at the cost of huge computational complexity and massive storage requirements. They possess the ability to learn and maneuver

billions of model parameters to discover non-linear patterns in the data. Despite their success, it is cumbersome to train them and deploy them on devices with limited resources such as smartphones and microcontroller-based devices. With this in mind, many researchers have developed model compression and acceleration techniques. Knowledge distillation is one of those techniques. The idea is to learn a small model, called the student with the help/guidance of a larger model, called the teacher [31]. Knowledge distillation leverages the reduced size of the student and at some level the performance of the teacher. Lightweight models offer real-time deployment in embedded devices. This quality is particularly recommendable for RTLS.

The theory behind the technique is established in [18], [32]. The authors present a vanilla version of the technique. But before explaining it, a little background is needed. In a classification problem, a neural network produces class probabilities by using a “softmax” output layer. This layer converts the logit computed for each class into a probability by comparing it with the other logit values. In this conversion exists a term called temperature “ T ”. “ T ” is normally set to 1, but using a larger number for “ T ” will produce a softer probability distribution over classes. The idea is to transfer knowledge from a big model, the teacher, with a high temperature in its softmax. Then the performance of a small distilled model, the student, is improved by using a weighted average of two different objective functions. The first objective function is the cross-entropy with the soft targets. This cross-entropy is computed using the same high temperature in the softmax of the distilled model as was used for generating the soft targets from the teacher model. The second objective function is the cross-entropy with the correct labels (hard targets) and a temperature of 1. For the weight of the objective functions, the authors found that the best results were generally obtained by using a considerably lower weight on the second objective function. A detail to consider is that the magnitudes of the gradients produced by the soft targets scale as $\frac{1}{T^2}$ it is important to multiply them by T^2 when using both hard and soft targets. In the ideal case, after applying knowledge distillation the output will be a smaller and capable model for the task at hand.

The teacher-student approach can also be used to transfer learning from one sensing modality to another. In [33], the authors augmented an IMU-based model with audio signals. Instead of training with only IMU data, an advanced audio-based teacher guided the training. The technique is capable of transferring rich context information from the audio-based model to the inertial-based model. There was an increment of 4.4% compared to the baseline (inertial model alone) for recognizing 23 activities of daily living. Once the student is trained, the human activity recognition (HAR) model only takes as input the motion data for inference. In [34], the authors use the same methodology of cross-modality transfer learning, which they call VAX. VAX means video/audio to X, where the video/audio is the teacher modality and X is a privacy-sensitive sensor. The motivation is that video/audio models are widely spread and trained with a large corpus of labeled training data. While privacy preserving modalities such

as mmWave Doppler radar and inertial measurement units (IMUs) do not readily generalize across environments and require significant in-situ training data. The author’s idea is to alleviate the lack of trained labeled data in privacy-aware sensor-based models with the guidance of the video/audio teacher. They achieved an improvement of 5% average accuracy compared to the student baseline (from 79% to 84%).

In this work, we propose to enhance the RSSI localization technique by distilling the knowledge of a multimodal teacher. The teacher fused RSSI signals with ultrasound coordinates to leverage multimodal information. Then, the student is guided by the multimodal teacher. Once the student is trained, the position of the user is inferred by a model whose inputs are only RSSI channels. The student is kept only with RSSI signal dependency, reducing cost, improving performance, and increasing the ubiquity and scalability of the solution.

III. METHOD

A. Hardware and Experimental Setting

The system consists of three parts and is distributed as shown in Fig. 1. The first part is a BLE tag attached to the volunteer’s upper arm. The second part is the RSSI receivers deployed around the factory modules. The reason for the selection of dedicated boards for the BLE tag and BLE receivers is that we recorded additional sensing modalities not available in smartphones, in addition to those presented in this paper. But ultimately, the scalability of our method depends on the possibility of using smartphones for the final application. And the third part is the ultrasound tracking system. The map’s setting aims to recognize in which module the worker is working (Positions 1-4).

The BLE tag consists of the development board Adafruit Feather Sense. The Adafruit board was attached to the participant’s upper arm with a smartphone Sport armband. Inside the feather sense, the BLE communication is handled by the Nordic nRF52832. It is a general-purpose multiprotocol SoC. The nRF52832 supports Bluetooth Low Energy, including the high-speed 2 Mbps feature. Moreover, Bluetooth mesh can be run concurrently with Bluetooth LE, enabling smartphones to provision, commission, configure and control mesh nodes. This means the device can be configured to interact with multiple nodes in a mesh network through a publish/subscribe paradigm. The RSSI resolution is 1 dB and the maximum transmission power is +4dBm (selected one).

The RSSI receivers were three MacBook-Pro deployed around the smart factory modules, running a Python RSSI reception script. The specifications of the receivers are in Table I. The position of the receivers is depicted in Fig. 1 in green color. The sampling rate of the RSSI signal is around 50 Hz.

The ultrasound localization system is the off-the-shelf system from Marvelmind Robotics. The version of the Marvelmind navigation package was 1.0.8 (2018). The system provides indoor positioning and navigation data with a precision of 2 cm. Ultrasound-based tracking has two elements; a hedge and the beacons. The hedge is then tracked by the ultrasound

TABLE I
BLE-RSSI RECEIVERS SPECIFICATIONS. ALL THE RECEIVERS WERE
MACBOOK PROS FROM DIFFERENT MODELS.

Model	OS	Position	Bluetooth Version
A2141-2019	Sonoma 14.3	1	5.0
A1502-2014	Big Sur 11.7.7	2	4.0
A1502-2015	Monterrey 12.7.2	3	4.0

beacons. Where the beacons' position is previously known. In our setting, the hedge is attached to the volunteer's back using a lifting belt. And, the ultrasound beacons were on the smart factory modules as shown in Fig. 1 in red color. The Marvelmind system used the information from the beacons and tracked the hedge in a three-coordinate space (X, Y, and Z). The sampling rate for the position is around 10 Hz.

B. Experimental Procedure

Twelve volunteers were recruited. They identify themselves as ten male and two female. Their age ranges from 23 to 59 years old (mean of 30.75). The height ranges from 160-184 centimeters (mean 178 cm). Only one of the participants was left-handed. The participants wore the RSSI-emitter (tag) on the left upper arm. The ultrasound hedge is placed on the volunteer's back using a lifting belt. At the beginning of the experiment, the sensors' data is synchronized in front of a video camera. The camera time is then used as a global clock to synchronize the data from the RSSI receivers and ultrasound beacons. The volunteers were asked to walk around the modules and simulate working activities on a factory floor. The activities performed by each volunteer were categorized as walking, touching the screen/buttons, opening the door, working inside the module, and closing the door. Each of these activities was performed on each module in each session. A total of five sessions per participant were recorded. In between every session, the hardware was removed from the wearer and a period of ten to twenty minutes rest was enforced. This makes the results accountable for the re-wearing of the system, which is typically expected in wearable devices. Each session lasted around 20 minutes on average. One participant performed two sessions one day and three sessions another. One volunteer only performs three sessions in total. And one volunteer has only four valid sessions. The participants with less than five sessions were only included in the training set (2 participants). For ten participants, the data is split into 4 sessions for training and 1 session for testing. A 5-fold cross-validation with a leave-one-session-out evaluation scheme is performed.³

C. BeSound Knowledge Distillation

The first step of the algorithm is the signal pre-processing. We obtained three BLE-RSSI signals from the BLE receivers and coordinates X, Y, and Z from the Marvelmind Ultrasound system. For the RSSI values a DC offset removal is performed by session (mean subtraction by channel). The subtraction of

the mean is intended to highlight signal variations rather than absolute values. The absolute values of the RSSI BLE signal vary even under static conditions, but the variations are comparatively stable. Therefore, a small variation indicates a stable position, and a high variation indicates relevant changes in the user's position. This helps the model to train faster, obtaining compatible results between sessions. This is followed by a 2-second resample window (100 samples at 50Hz). Then, a Butterworth low pass filter of 3 Hz was used to remove the jitter on the RSSI signal. For the case of ultrasound-based coordinates a resample to 2 seconds window is performed. This was followed by a Butterworth low pass filter of 10 Hz to remove the ringing from the resampling procedure and keep the pattern of the original sampling rate of 10 Hz. The resampling to 50Hz is for synchronization purposes with the video-ground truth with 50 frames per second. The data is manually annotated according to the position of the participants in the recorded video. A sliding window of 2 seconds with a step size of 0.5 seconds is employed. After pre-processing, the data set structure is six channels with a window size of 2 seconds and 25% overlapping. The ground truth of the worker's position is extracted from the recorded videos manually. The categories were four positions related to the work modules depicted in Fig. 1. The next step in our approach is to train the teacher for the knowledge distillation to the student. The TensorFlow 2.15.0 version is used to train the networks. The evaluation scheme is defined as a 5-fold cross-validation with leave-one-session out. The training of teacher and student baselines ran for 100 epochs with early stopping (patience 20) and restoring the best weight option enabled to avoid overfitting. The training for the distilled model ran for 10 epochs. And, a 64-batch size was selected.

The Teacher: In this work, we have trained a multimodal and multipositional teacher model. RSSI-BLE signals from the three receivers and the X, Y, and Z coordinates of the ultrasound localization system were fused. This fusion is considered multimodal and multipositional due to the sensor's distribution on the experimental map in Fig. 1. The structure of the teacher model is depicted in Fig. 2 **Left**. The input layer consists of six channels, three BLE-RSSI, and three ultrasound position coordinates. These six channels are fed to two separate neural networks (NNs). One of the neural networks is focused on extracting features of the cross-channel interaction between the modalities and positions. The details of the NN for cross-channel interaction extraction are in Table II. And the second neural network aims to extract the causality of the multimodal time series input. Both networks are then concatenated and fed into a classifier layer (Linear layer), followed by an output layer with the softmax activation function. The causality extractor network is based on Long-Short-Term Memory (LSTM) [36], and the details are in Table III. The combination of the two concatenate NNs can capture spatial and temporal information simultaneously, thus effectively solving complex time series problems. With this structure, the teacher has 4594 parameters (79.75 KB). The optimizer is Adam with a learning rate of 0.001. The loss

³All participants signed an agreement following the policies of the university's committee for protecting human subjects and following the Declaration of Helsinki [35].




Fig. 2. BeSound Knowledge Distillation Approach. **Left** The Multimodal and Multipositional Teacher Consists of Two Concatenate Networks; One For Cross-Channel Interaction Feature Extraction and One Network for Causality Extraction (LSTM-Based). **Right** The Student Consists of a Multipositional Feature Extractor and a Classifier. The Distillation is Applied at the Logit Level to Improve the Position Estimation of the Student.

TABLE II
IMPLEMENTATION DETAILS OF THE NEURAL NETWORK FOR
CROSS-CHANNEL INTERACTION OF THE TEACHER MODEL

Layers	Configuration
Dense	Size 10 Activation Function ReLu
Conv1D-1	Filters 40 Kernel 10 Padding "Same" Activation Function ReLu MaxPooling Size 10 Dropout 0.5
Conv1D-2	Filters 40 Kernel 10 Padding "Same" Activation Function ReLu MaxPooling Size 10 GlobalAvgPooling
Dense	Size 10 Activation Function ReLu

TABLE III
IMPLEMENTATION DETAILS OF THE NEURAL NETWORK FOR CAUSALITY
EXTRACTION OF THE TEACHER MODEL

Layers	Configuration
LSTM	Units 16 Return Sequences True Flatten
Dense	Size 10 Activation Function ReLu

function is categorical cross-entropy and the metric to monitor is accuracy.

The Target Model: The idea is to explore the performance enhancement of small and simple networks that can be deployed in wearable and embedded devices with a reduced impact on power consumption and memory. The structure of the student/target NN is depicted in Fig. 2 **Right**. The input layer consists of three BLE-RSSI channels. The network combines a multipositional feature extraction and a classifier with two linear layers. The NN is a CNN-based model. Three sizes of student models were evaluated. The sizes are defined by the number of filters and kernels of the CNN layers. The number of parameters for each student is 754, 1824, and 4594, respectively. Which is around 27.07, 11.19, and 4.44 times smaller than the teacher with 20416 parameters, respectively. The details of the networks are in Table IV The optimizer is Adadelata with a learning rate of 0.9. The loss

TABLE IV
IMPLEMENTATION DETAILS OF THE NEURAL NETWORK OF THE
STUDENT/TARGET MODEL

Layers	Student/Target 1	Student/Target 2	Student/Target 3
	Filters 5 Kernel 3 Padding "Same" ReLu MaxPool 5 Dropout 0.2	Filters 10 Kernel 5 Padding "Same" ReLu MaxPool 5 Dropout 0.2	Filters 20 Kernel 5 Padding "Same" ReLu MaxPool 5 Dropout 0.2
Conv2D-2	Filters 5 Kernel 3 Padding "Same" ReLu MaxPool 5 Flatten	Filters 10 Kernel 5 Padding "Same" ReLu MaxPool 5 Flatten	Filters 20 Kernel 5 Padding "Same" ReLu MaxPool 5 Flatten
Dense	Size 10 ReLu Dropout 0.1	Size 10 ReLu Dropout 0.1	Size 10 ReLu Dropout 0.1

function is categorical cross-entropy and the metric to monitor is accuracy.

Distilled Student: In this work, the vanilla version of knowledge distillation is applied [18]. The optimizer is Adam [37] with a learning rate of 0.001. The distilled students were trained with the goal of minimizing the student loss and the distillation loss. The equation for the total loss is in Eq. (1). Where α (set to 0.1) is the weight for the objective function, sl is the student loss and dl is the distillation loss.

$$loss = \alpha \times sl + (1 - \alpha) \times dl \quad (1)$$

The student loss is defined as sparse categorical cross-entropy loss with logit enabled. The distillation loss is defined in Eq. (2). Where T is the distillation temperature, set to 10. KLD is the Kullback-Leibler divergence loss (KLD). KLD is a non-symmetric metric that measures the relative entropy or the difference in information represented by two distributions. It measures the distance between two data distributions to show how far they are from each other. This metric is typically used in knowledge distillation tasks [38].

$$dl = KLD(V_t, V_p) * T^2 \quad (2)$$

The true values of KLD (V_t) are defined as the softmax function in Eq. (3). Where $\vec{z}$ is the input vector defined as $\frac{T_p}{T}$, where T_p is the teacher prediction and T is the distillation




Fig. 3. Multimodal and Multipositional Teacher Average Results with 5Fold-Cross Validation with Leave-Session-Out Evaluation Scheme. **Left** Confusion Matrix Result for a Window Size of Two Seconds with F1-Score of 79.85%. **Right** Confusion Matrix Result for a Window Size of Ten Seconds with F1-Score of 84.41%.

temperature. e^{z_i} is the exponential of input vector, e^{z_j} is the exponential of the output vector and K is the number of categories. The predicted values of KLD (V_p) are defined as the softmax function in Eq. (3). Where $\vec{z}$ is defined as $\frac{S_p}{T}$, where S_p is the student prediction and T is the distillation temperature.

$$\sigma(\vec{z}) = \frac{e^{z_i}}{\sum_{j=1}^K e^{z_j}} \quad \text{for } i = 1, 2, 3, \dots, K \quad (3)$$

IV. RESULT AND DISCUSSION

Fig. 3 presents the results of the trained teacher. An F1-score of 79.85% and 84.41% is obtained for the two-second and ten-second window recognition, respectively. The highest confusion pair is Position 3 and Position 4. This might be because of our experimental setting map (see Fig. 1). For the case of BLE-RSSI signals, Position 3 and Position 4 are covered by only one RSSI receiver. Unlike Position 1 and Position 3, both have their own RSSI receiver. In the case of the ultrasound beacons, we observed that volunteers to moved between Position 3 and Position 4 tended to select the path behind the factory modules, where the ultrasound antennas were not pointing. When a 10-second recognition window is selected, the results improve by 4.56%, indicating that the detection modalities need time to stabilize after the participant moves between modules.

Fig. 4 (a) depicts the confusion matrices for the target models' baselines to be later compared with the distilled student models. The differences between the students/targets are the number of filters and kernels in the CNN layers. The first student has only 754 parameters, 27 times smaller than the teacher (20416 parameters). And, the biggest student has 4594 parameters, which is 4 times smaller than the teacher model. All the target models' baselines present the highest misclassification when estimating Position 1. The Position 1 as shown in Fig. 1 is 3 times bigger than the rest of the factory modules. BLE-RSSI localization is not suitable for long-range positioning. BLE technology has been developed for short-range communications up to 10 meters. The 10-meter range is without considering interference from surrounding obsta-


cles/people. Thus, determining Position 1 can be challenging with only BLE-RSSI signals as input.

Fig. 4 (b) presents the confusion matrices of the distilled students. For the case of Distilled Student 1 an F1-score of 67.08 % is achieved. This result indicates an 8.98% performance increase compared to the target model 1 baseline with 58.10%. The student is 27 times smaller than the teacher and it can achieve an F1-score difference of 12.77% compared to the teacher (79.85% F1-score). The medium-size model, Distilled Student 2, has a performance of 12.06% F1-score lower than the teacher and 8.31% higher than the target model 2 baseline. The Distilled Student 3 outperformed the target model baseline by 11.79% F1-score, and the difference between the student and the teacher was reduced to 8.89%. This model is 4 times smaller than the teacher. The recall of Position 1 is highly increased for all the distilled models compared to the target model baselines (see Fig. 4). Increasing the number of parameters of the student from 754 to 4594 (16% increment) has a 3.88% increase in the F1-score.


Overall, our results show the potential of cross-modal knowledge distillation for localization systems. All distilled student models outperformed the target model baselines. This increase in performance indicates that the multimodal and multipositional teacher effectively guides the student. With this method, the model is compressed and improved. Moreover, in the inference, only the BLE-RSSI signals are the inputs. This means that our solution retains the ubiquity and low-cost advantages of BLE-RSSI and, at the same time, improves performance.

On the other hand, our solution has several **limitations** and possibilities for improvement. This is not a complete list but they are relevant limitations of our work.

- **Detection Radio:** The detection radio is dependent on the number of receivers for the wearable BLE RSSI. In this case, three RSSI receiver devices were deployed (one per module) in the environment. In the future, it will be interesting to remove the dependency on sensors on the environment. This could be achieved by distilling cross-modal knowledge to a learner whose inference is only based on IMU data.
- **Experimental Setting:** The experimental setting map in Fig. 1 shows a minimum distance between the factory modules of 2 meters. The method will increase in relevance if it is possible to capture distance with an accuracy of centimeters. The experiment is performed by one volunteer at a time with a maximum of 8 people around and a minimum of 2 people around. Hence, additional challenges might arise if the positions are estimated in crowded spaces.
- **Temperature Value and Weight for Objective Function:** We have selected a temperature value of 10 and a weight of 0.1 for the objective function of the knowledge distillation. This selection follows typical values of those parameters for the vanilla version of distillation. But, it might be valuable to test different values and analyze the performance results.



(a) Unimodal and Multipositional Target Model Baselines Average Results with 5Fold-Cross Validation with Leave-One-Session Out Evaluation Scheme. **Target Model 1 (754 Parameters)** Results for a Window Size of Two and Ten Seconds with F1-Score of 58.10% and 60.19%, respectively. **Target Model 2 (1824 Parameters)** Results for a Window Size of Two and Ten Seconds with F1-Score of 59.48% and 61.91%, respectively. **Target Model 3 (4594 Parameters)** Results for a Window Size of Two and Ten Seconds with F1-Score of 59.10% and 62.31%, respectively.



(b) Distilled Student Models Average Results with 5Fold-Cross Validation with Leave-One-Session Out Evaluation Scheme. **Distilled Student 1** Results for a Window Size of Two and Ten Seconds with F1-Score of 67.08% and 72.84%, respectively. **Distilled Student 2** Results for a Window Size of Two and Ten Seconds with F1-Score of 67.79% and 73.17%, respectively. **Distilled Student 3** Results for a Window Size of Two and Ten Seconds with F1-Score of 70.96% and 76.04%, respectively.

Fig. 4. Target Models and Distilled Students Models Average Results with 5Fold-Cross Validation with Leave-One-Session Out Evaluation Scheme.

- **Kullback-Leibler Divergence Loss (KLD):** We employ the KLD loss because this loss has achieved considerable success by controlling the “soft” targets via the temperature scaling parameter. In [38], the authors proved that the MSE loss outperforms the KLD loss, explained by the difference in the penultimate layer representations between the two losses.
- **Knowledge Distillation at Logit Level:** The method we have used is based on the distillation loss feedback in the output layer. This can be extended to the feature layer in the student network. Moreover, there exist other methods such as self-distillation [39], learning from multiple teachers [40], and combining knowledge distillation with data-free compression [41], among others.
- **Smartphone Antennas Variability:** In our experiment design we purposely employed Apple products for the BLE receivers due to their stability in receiving the BLE RSSI. However, we have used different versions of MacBooks to add variability in the receiving antennas, which is to be expected in the real-world scenario where there is a huge variety of smartphone vendors. The variability included in the experiment may not be sufficient to represent the variety of smartphone antenna designs, so in the future, it is pertinent to quantify the impact of the hardware on the signal reception.

V. CONCLUSION

In this paper, we present a framework for improving a localization system based on BLE-RSSI signals. We have in-

troduced multimodal and multipositional BLE-RSSI and ultrasonic coordinates information using a teacher-student framework. The BLE-RSSI student model benefits from teacher guidance during training, while relying solely on BLE-RSSI data for inference. This approach retains the advantages of ubiquity, low cost, and scalability of the BLE-RSSI positioning method. Based on a study with 12 participants in which they simulated working and moving between modules of the smart factory test bed, we show the extent to which our framework improved performance. We have evaluated the approach with three student/target models. The students were designed with the same neural network structure but with a different number of parameters. Experimental results show that with the smallest student with 754 parameters (27 times smaller than the teacher), the F1-score is 8.98% higher than the target model baseline. Moreover, for the biggest student with 4594 parameters (4 times smaller than the teacher), the F1 score is 11.79% higher than the target model baseline. These results were achieved by training the distilled students for 10 epochs. In contrast to the target model’s baselines trained for 100 epochs. This makes our approach efficient and sustainable with fewer epochs per training and higher performance. We also discussed a list of relevant limitations of our approach with future solutions.

ACKNOWLEDGMENT

The research reported in this paper was partially supported by the German Federal Ministry of Education and Research (BMBF) in the projects SocialWear (01IW20002) and Vid-

GenSense (01IW21003). We would also like to thank the volunteers and the student help of Luis Sanchez.

REFERENCES

- [1] F. B. Islam, J.-M. Lee, and D.-S. Kim, "Smart factory floor safety monitoring using uwb sensor," *IET Science, Measurement & Technology*, vol. 16, no. 7, pp. 412–425, 2022.
- [2] S. Yang, L. Ma, S. Jia, and D. Qin, "An improved vision-based indoor positioning method," *IEEE Access*, vol. 8, pp. 26 941–26 949, 2020.
- [3] D. Maeda, Y. Watanabe, G. Joly, and S. Gaur, "Fusion of vision deep-learning person detection and radio beacon ranging for identifying and locating factory workers," *IEEE Sensors Letters*, 2023.
- [4] D. Khan, Z. Cheng, H. Uchiyama, S. Ali, M. Asshad, and K. Kiyokawa, "Recent advances in vision-based indoor navigation: A systematic literature review," *Computers & Graphics*, vol. 104, pp. 24–45, 2022.
- [5] R. Liu, C. Yuen, T.-N. Do, and U.-X. Tan, "Fusing similarity-based sequence and dead reckoning for indoor positioning without training," *IEEE Sensors Journal*, vol. 17, no. 13, pp. 4197–4207, 2017.
- [6] J. Lee, Y. Park, D. Kim, M. Choi, T. Goh, and S.-W. Kim, "An efficient localization method using rfid tag floor localization and dead reckoning," in *2012 12th International Conference on Control, Automation and Systems*. IEEE, 2012, pp. 1452–1456.
- [7] C.-C. Tsai, "A localization system of a mobile robot by fusing dead-reckoning and ultrasonic measurements," *IEEE Transactions on Instrumentation and Measurement*, vol. 47, no. 5, pp. 1399–1404, 1998.
- [8] A. Velten, T. Willwacher, O. Gupta, A. Veeraraghavan, M. G. Bawendi, and R. Raskar, "Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging," *Nature communications*, vol. 3, no. 1, p. 745, 2012.
- [9] A. Velten, D. Wu, A. Jarabo, B. Masia, C. Barsi, C. Joshi, E. Lawson, M. Bawendi, D. Gutierrez, and R. Raskar, "Femto-photography: capturing and visualizing the propagation of light," *ACM Transactions on Graphics (ToG)*, vol. 32, no. 4, pp. 1–8, 2013.
- [10] D. B. Lindell, G. Wetzstein, and M. O'Toole, "Wave-based non-line-of-sight imaging using fast fk migration," *ACM Transactions on Graphics (ToG)*, vol. 38, no. 4, pp. 1–13, 2019.
- [11] S. Yue, H. He, P. Cao, K. Zha, M. Koizumi, and D. Katabi, "Cornerradar: Rf-based indoor localization around corners," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 6, no. 1, pp. 1–24, 2022.
- [12] U. M. Qureshi, Z. Umair, and G. P. Hancke, "Indoor localization using wireless fidelity (wifi) and bluetooth low energy (ble) signals," in *2019 IEEE 28th International Symposium on Industrial Electronics (ISIE)*. IEEE, 2019, pp. 2232–2237.
- [13] Z. Ma, S. Poslad, J. Bigham, X. Zhang, and L. Men, "A ble rssi ranking based indoor positioning system for generic smartphones," in *2017 Wireless Telecommunications Symposium (WTS)*. IEEE, 2017, pp. 1–8.
- [14] A. Lindemann, B. Schnor, J. Sohre, and P. Vogel, "Indoor positioning: A comparison of wifi and bluetooth low energy for region monitoring," in *HEALTHINF*, 2016, pp. 314–321.
- [15] A. A. Morgan, G. S. B. Humaid, and A. I. Moustafa, "Using bluetooth low energy for positioning systems within overcrowded environments: A real in-field evaluation," *Computers & Electrical Engineering*, vol. 93, p. 107314, 2021.
- [16] F. Ijaz, H. K. Yang, A. W. Ahmad, and C. Lee, "Indoor positioning: A review of indoor ultrasonic positioning systems," in *2013 15th International Conference on Advanced Communications Technology (ICACT)*. IEEE, 2013, pp. 1146–1150.
- [17] S. Roa-Prada, H. A. Scarton, G. J. Saulnier, D. A. Shoudy, J. D. Ashdown, P. K. Das, and A. J. Gavens, "An ultrasonic through-wall communication (utwc) system model," *Journal of Vibration and Acoustics*, vol. 135, no. 1, p. 011004, 2013.
- [18] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.
- [19] F. Milano, H. da Rocha, M. Laracca, L. Ferrigno, A. Espírito Santo, J. Salvado, and V. Paciello, "Ble-based indoor localization: Analysis of some solutions for performance improvement," *Sensors*, vol. 24, no. 2, p. 376, 2024.
- [20] G. Bahle, V. Fortes Rey, S. Bian, H. Bello, and P. Lukowicz, "Using privacy respecting sound analysis to improve bluetooth based proximity detection for covid-19 exposure tracing and social distancing," *Sensors*, vol. 21, no. 16, p. 5604, 2021.
- [21] S. Iakimenko, A. Tsislav, E. Naumova, I. Bessonova, A. Novikov, and N. Seltsova, "Bluetooth 5.1 indoor localization study based on angle-of-arrival method," *Telecommunications and Radio Engineering*, vol. 81, no. 3, 2022.
- [22] A. K. Mustafa and E. R. Sykes, "A high fidelity indoor navigation system using angle of arrival and angle of departure," in *Proc. Work Prog. Papers 13th Int. Conf. Indoor Positioning Indoor Navigation*, 2023.
- [23] M. Cominelli, P. Patras, and F. Gringoli, "Dead on arrival: An empirical study of the bluetooth 5.1 positioning system," in *Proceedings of the 13th international workshop on wireless network testbeds, experimental evaluation & characterization*, 2019, pp. 13–20.
- [24] D. Giovanelli and E. Farella, "Rssi or time-of-flight for bluetooth low energy based localization? an experimental evaluation," in *2018 11th IFIP wireless and mobile networking conference (WMNC)*. IEEE, 2018, pp. 1–8.
- [25] A. Comuniello, A. De Angelis, A. Moschitta, and P. Carbone, "Using bluetooth low energy technology to perform tof-based positioning," *Electronics*, vol. 11, no. 1, p. 111, 2021.
- [26] D. Giovanelli, E. Farella, D. Fontanelli, and D. Macii, "Bluetooth-based indoor positioning through tof and rssi data fusion," in *2018 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, 2018, pp. 1–8.
- [27] D. Zhu and J. Yan, "A deep learning based bluetooth indoor localization algorithm by rssi and aoa feature fusion," in *2022 International Conference on Computer, Information and Telecommunication Systems (CITS)*. IEEE, 2022, pp. 1–6.
- [28] A. Koutiris, T. Siozos, Y. Kopsinis, A. Pikrakis, T. Merk, M. Mahlig, S. Papaharalabos, and P. Karlsson, "Deep learning-based indoor localization using multi-view ble signal," *Sensors*, vol. 22, no. 7, p. 2759, 2022.
- [29] S. Bian, P. Hevesi, L. Christensen, and P. Lukowicz, "Induced magnetic field-based indoor positioning system for underwater environments," *Sensors*, vol. 21, no. 6, p. 2218, 2021.
- [30] R. Amsters, E. Demeester, N. Stevens, Q. Lauwers, and P. Slaets, "Evaluation of low-cost/high-accuracy indoor positioning systems," in *Proceedings of the 2019 International Conference on Advances in Sensors, Actuators, Metering and Sensing (ALLSENSORS)*, Athens, Greece, 2019, pp. 24–28.
- [31] J. Gou, B. Yu, S. J. Maybank, and D. Tao, "Knowledge distillation: A survey," *International Journal of Computer Vision*, vol. 129, pp. 1789–1819, 2021.
- [32] C. Bucilua, R. Caruana, and A. Niculescu-Mizil, "Model compression," in *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2006, pp. 535–541.
- [33] D. Liang, G. Li, R. Adaimi, R. Marculescu, and E. Thomaz, "Audioimu: Enhancing inertial sensing-based activity recognition with acoustic models," in *Proceedings of the 2022 ACM International Symposium on Wearable Computers*, 2022, pp. 44–48.
- [34] P. Patidar, M. Goel, and Y. Agarwal, "Vax: Using existing video and audio-based activity recognition models to bootstrap privacy-sensitive sensors," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 7, no. 3, pp. 1–24, 2023.
- [35] D. A. Shephard, "The 1975 declaration of helsinki and consent," *Canadian Medical Association Journal*, vol. 115, no. 12, p. 1191, 1976.
- [36] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [37] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [38] T. Kim, J. Oh, N. Kim, S. Cho, and S.-Y. Yun, "Comparing kullback-leibler divergence and mean squared error loss in knowledge distillation," *arXiv preprint arXiv:2105.08919*, 2021.
- [39] L. Zhang, J. Song, A. Gao, J. Chen, C. Bao, and K. Ma, "Be your own teacher: Improve the performance of convolutional neural networks via self distillation," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 3713–3722.
- [40] S. You, C. Xu, C. Xu, and D. Tao, "Learning from multiple teacher networks," in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2017, pp. 1285–1294.
- [41] R. G. Lopes, S. Fenu, and T. Starner, "Data-free knowledge distillation for deep neural networks," *arXiv preprint arXiv:1710.07535*, 2017.